Calculating Total Cost of Ownership on Intrusion Prevention Technology

Sponsored by Sourcefire

Calculating Total Cost of Ownership on Intrusion Prevention Technology

February 2014

A SANS Analyst Product Review

Written by Eugene E. Schultz, Ph.D.

Updated by J. Michael Butler

Advisors: J. Michael Butler & Dave Shackleford

Value Proposition *PAGE 2*

Methodology *PAGE 4*

TCO Exercises Favor Automated Management *PAGE 5*

Reduced Exposure = Cost Avoidance *PAGE 9*

Executive Summary

Advanced attacks, malware and evasion techniques are challenging intrusion prevention systems (IPSes) to be smarter, faster and more accurate. The terms *advanced IPS* or *next-gen IPS* (NGIPS) may involve a firewall and IPS appliance working as one. These systems work together to help IPSes make more informed decisions and detect and block undesirable events before they have a negative impact on downstream systems.

With the most accurate information available, NGIPSes are able to intelligently *intervene*, rather than simply send alerts. If the NGIPS can accurately detect and terminate a disruptive and potentially costly security-related incident, it can save an organization what could be a sizeable expenditure related to remediation, system interruption, data loss and possible loss of reputation.

With advanced correlation and automation, there are many areas in which NGIPSes can save organizations time and money—particularly in correlating the applicability of the perceived event to the organization's actual vulnerability posture.

This paper, while not scientific, attempts to calculate the value of specific automation features in NGIPSes with which organizations can achieve savings in total cost of ownership (TCO). The paper is designed to help organizations expand this TCO concept to determine realistic savings they could potentially achieve in their environments as NGIPS tools embed more automated features.

Intrusion Prevention Requirements

Next-gen intrusion prevention systems (NGIPSes) must detect anomalies within both inbound and outbound packets with more speed and accuracy. An NGIPS must be able to interface with other security tools, such as decryption, whitelisting, firewalls, analytics/intelligence platforms, security information and event management systems (SIEMs) and other dashboard devices for correlation and analysis. Major features and functions of NGIPSes include the following capabilities:

- Accept regular updates on suspicious patterns, applications and malware.
- Inspect traffic down to the data level without impeding legitimate traffic.
- Block non-allowed network traffic, applications, incoming services and other requests to hosts in accordance with organizational policy.
- Support both passive detection and active blocking based on policy.
- Look into encrypted packets (usually through additional decryption technologies).
- Collect accurate data for other analytics, SIEM and firewall systems.
- Collect and preserve data that can be used easily for analysis and forensics purposes.
- Calculate and display high-level data, such as in a dashboard.
- Failover safely if something interrupts the operation of the NGIPS.

Value Proposition

The current consensus among information security professionals is that ROI is difficult to achieve in the realm of information security. ROI is typically calculated in connection with evaluating the success of activities and methods designed to earn financial profit for an organization. Because it is not a revenue generation engine for an organization, information security efforts focus on striving to reduce losses by percentages or amounts set by executive-level management. So, rather than trying to prove ROI in connection with their information security efforts, organizations typically are attempting to achieve reductions in TCO related to managing their information security practices.

Advances in information security technology have resulted in products that are less expensive to purchase and require less labor to install and maintain, all while delivering more critical functionality than ever before. The result is savings in terms of time and monetary cost when compared to more traditional controls, many of which may involve manual procedures. For example, time to respond to incidents would be a category in which TCO could be improved by automating the process of looking up associated end users with IP addresses and network segments that have been attacked.

Although there are many areas in which an NGIPS can earn back its value, we've determined four TCO savings areas in which security automation may have the greatest effect:

1. **Automated tuning.** Time involved in initial and ongoing tuning of IPSes can be measured. IPSes need to be tuned *before* they start working. They must work with network monitoring systems to know what machines are on the network and the vulnerabilities associated with those machines and systems. Once an IPS is running, security personnel need to tune its configuration parameters continuously so that it is aware of which machines are added to and removed from networks, the vulnerabilities associated with those machines and so on. With an NGIPS, security policy recommendations can be automated. Automated tuning through network monitoring mechanisms that identify malicious and normal behavior and then adjust rules accordingly reduces TCO compared to completing such tasks manually.

2. **Impact assessment.** False positives—or alerts that are actually nonevents—consume huge amounts of resources. An intelligent NGIPS will work in conjunction with its own asset map and/or an external asset management system to determine whether an alert may have high impact or whether the alert is a nonevent because the network has no target for that exploit. When potentially adverse events occur, the NGIPS must make a judgment concerning their impact early in the incident response process. For example, an attack against the remote procedure call (RPC) in a Windows system will not succeed if the target is a Linux system. This event would be deemed a low-impact event requiring no intervention; however, all events should ultimately be reviewed because they may create other problems on the network if left unchecked.

   Higher impact events, such as a SQL injection attack that applies to your version and patch level of Apache server, are often called *actionable* events. Such an event would, then, generate an alert and response.

3. **Linking individual users with events.** Because most infections begin with endpoints, identifying the user involved and being able to talk to that person—and being able to cut off his or her access quickly—is imperative. Also knowing the source of the infection will help speed up the determination as to where that infection is attempting to spread. Many organizations still look up user directories manually to locate and identify users associated with affected nodes. This process can be time-consuming in complex enterprises. Automatically correlating assessed actual events to the activity of specific users (usually derived from user directories and network discovery mechanisms) can result in locating each user within seconds rather than hours.

4. **Loss prevention/cost avoidance.** The *prevention* part of NGIPS is, of course, the most critical cost-saving function that the NGIPS can provide for an organization. Once the intrusion starts to spread, detection and remediation costs rise—as does the risk of data loss. Data breaches due to a malicious attack cost organizations $275 per record to remediate in 2012, according to the 2013 Cost of a Data Breach Report by Ponemon Institute.[1] Applying this estimated cost to the recent Target data breach, now pegged at 70 million records,[2] would project a total hit to Target's bottom line of $19.25 billion. Although saving that expense is not a TCO or ROI element, per se, we must consider our organization's capability to survive the material impact of data loss and how much we are willing to invest in order to avoid such losses in regard to sensitive data we store in our systems.

[1] www.bankinfosecurity.com/interviews/data-breach-i-1953/op-1, graph on page 1

[2] https://corporate.target.com/about/shopping-experience/payment-card-issue-FAQ#q5888

Methodology

This study is based on real-life experience, applicable outside research and events, and a user case study; however, we do not claim it to be a scientific study.

In the first half of this study, we derive TCO calculations for a sample enterprise environment based on the cost of man-hours involved with manually *managing* three areas of IPS (IPS tuning, accurate impact assessment and linking users to actual security events for quicker forensics/remediation). Organizations wishing to emulate our process can do so by scaling their organizational size and creating similar manpower equations.

Sample Network

In the sample organization, we created a network for a larger company expected to have 7,500 users and 10,000 nodes distributed among 5 perimeter locations and 16 internal network points. Each perimeter location has its own (inline) IPS positioned behind the network firewall. Each internal location has an IPS configured in passive alerting mode.

Calculations

Calculations performed in this study compare the difference in time and manpower using manual versus automated methodologies in three cost-reduction areas for IPS: tuning, impact assessment and linking users to events. In this study, reduction in labor hours is calculated using the rate of $75/hour, a rate set by NSS Labs for the labor cost of IPS tuning.[3] To avoid overcomplicating our calculations, we consider all man-hour rates at $75, even though some may be lower or higher due to specific skill sets required to respond to incidents and other variables. Each organization needs to set this rate to its own pay scale to be able to determine its own TCO for each area covered in this report.

We discuss cost avoidance in the "Reduced Exposure = Cost Avoidance" section of the paper.

[3] www.darkreading.com/security/vulnerabilities/showArticle.jhtml?articleID=222001334

TCO Exercises Favor Automated Management

Using the criteria set in the "Methodology" section, we calculated an overall savings to be achieved through three areas of IPS automation: tuning, impact assessment and linking users to actual security events.

Automated Tuning

Annual $avings:
$39,720

An IPS needs to be tuned regularly to maximize the probability that events that constitute potential or actual attacks are detected and responded to, while at the same time reducing false alarms to negligible levels. Tuning requires technical staff to have enough knowledge of their traffic and logs to be able to evaluate the results of the IPS system to validate accuracy. This activity is not something that is completed immediately after an IPS is installed and initially configured; instead, it must occur continuously as the IPS becomes familiar with the network and administrators get familiar with the IPS.

The greatest amount of time spent tuning occurs during the first four weeks after an IPS is put into operation, with ongoing tuning being intermittent. For our 7,500-node environment, we estimate that manual tuning of the first IPS would require a minimum of 16 hours of labor on the part of technical staff over an initial period of one month. So, for the initial startup of this automation, our calculation looks like this:

Cost of initial IPS setup/tuning and policy creation, initial month
= 16 hours x $75/hour
= **$1,200**

The cost of tuning each of the additional IPSes in our hypothetical network would, again, be reduced because of the previously discussed learning factor. Assuming that the time required to tune each additional SIEM tool after the first would be 30 percent of initial setup and policy creation, we would calculate the time needed to tune our network to be 16 hours multiplied by 30 percent, or 4.8 man-hours for the remaining devices.

Not all of these devices have separate policies requiring separate tuning. Let's say, because of geography and the nature of their business, each of the five external network branches has unique elements in its IPS policies. Thus, five perimeter IPS policies would need to be tuned in an ongoing basis. Further, let's say that internal IPS policies are concentric and represent four separate enforcement policies (web server, data center and so on).

The first month would also include initial tuning of the eight remaining policies at the reduced 4.8 hours per policy. Given these variables, the first-month cost of tuning the eight remaining policies plus the initial policy tuning would be:

8 additional detection policies x 4.8 hours = 38.4 hours
38.4 hours x $75/hour = $2,880
$2,880 + $1,200 (initial policy tuning) = **$4,080** for tuning all policies during first month

Each of the policies will need to be tuned on a monthly basis. So, the time needed to manually tune nine separate IPS policies per month at 4.8 hours each would be:

9 policies to tune x 4.8 hpm (hours per month) = 43.2 hpm
43.2 hpm (for tuning 9 distinct IPS policies) x $75/hour = $3,240
$3,240 x 11 months following initial tuning = **$35,640** for 11 months of IPS policy tuning

To calculate the TCO estimate for 12 months of manually tuning the nine IPS policies for all IPS devices, add the initial cost to set up the policies and the costs associated with the remaining 11 months:

$4,080 for initial month + $35,640 for 11 remaining months = **$39,720** per year

These costs could be mostly eliminated if the IPS devices could automatically tune themselves, although some follow-up by technical staff would still be required.

Automated Impact Assessment

Annual Savings:
$108,000

Impact assessment means correlating a variety of information about an attack, the target(s) of the attack and the effect of the attack on an organization's processes and assets to know which events require action. In our sample organization—and without a centralized, automated operations center to analyze each event—IT staff could easily be drowned in hundreds of thousands of alerts that may or may not impact their network.

The amount of time required to assess the impact of these alerts depends on the scope and magnitude of the incident and often requires the input of a team of stakeholders such as the information security manager, the head of risk management, a legal representative, a human relations manager and others whose hourly rate exceeds $75/hour. But to keep things simple, we will calculate labor costs at the $75/hour rate when we get to our equation.

In our organization, we can presume that IPS sensors are triggering what, conservatively, could be hundreds of thousands of events per month. By intuition and human knowledge of the network, security analysts can tune out a large percentage of those. However, because networks are constantly changing and new threats emerge daily, the analysts can't possibly know everything about their systems, networks and traffic patterns. So conservatively, the security analysts would still be distracted by thousands of raw IPS events on a daily basis. Based on an interview with an IT security manager from our case study organization comprised of 20,000 nodes and 7,500 users (see Appendix A), the security staff spent approximately 160 man-hours per month to assess the impact of raw IPS security events.

Because our sample organization has half the number of nodes as our case study organization—but the same number of users—we can reduce that figure to 75 percent of the case study hours, or 120 hpm. If handled mostly manually, the costs for assessing impact, then, calculates this way:

120 hpm x $75/hour = $9,000 per month
$9,000 per month x 12 months = **$108,000** per year

The cost of filtering through large quantities of raw IPS events to uncover which events are applicable can be virtually eliminated if the NGIPS can automatically assess the impact of raw IPS events. This can be accomplished by the NGIPS's management console correlating threats against host/endpoint intelligence collected by the IPS and known vulnerabilities associated with operating systems and applications related to attacks—although some follow-up by technical staff may still be required.

Linking Individual Users with IPS Events

Annual $avings:
$37,125

Not surprisingly, linking users to IPS events is a large part of the expense associated with an IPS because, in most cases, DHCP is used to assign IP addresses to end-user devices. Because IP addresses can change frequently outside the DMZ, certain hosts are nearly impossible to identify with an IP address alone.

To approximate the costs recoverable through automated user identification, let's refer again to our case study, in which the company realized a 99 percent reduction of actionable events with intelligent NGIPS filtering, leaving them with 200 actionable events per month. Because our organization contains 10,000 fewer nodes but the same number of employees as our case study, we can assume more than half of this number of actionable events would be occurring on our sample network. So, let's say we're looking at 125 actionable events per month in our sample network.

Let's further estimate that two-thirds (or 67%) of those events represent servers with static IP addresses (e.g., DMZ, data centers) and one-third (33%) of the events involve end-user devices with IP addresses assigned through DHCP. (End-user devices can also be the source of an attack within an organization, whether linked to malicious users or users unknowingly propagating malware.)

Without an automated capability to correlate Active Directory or Lightweight Directory Access Protocol (LDAP) usernames with IP addresses, security analysts are left to sift through log files manually. This process can consume an hour or longer per inquiry. For purposes of this TCO analysis, let's assume one hour per inquiry.

With these assumptions in mind, TCO benefits can be calculated for our sample enterprise as follows:

125 actionable events x .33 (users with DHCP) = 41.25 manual lookup events
41.25 lookups a month x 1 hour at $75/hour = $3,093.75
$3,093.75 per month x 12 months = **$37,125** per year

So, when totaling the amount of money spent manually correlating usernames associated with actionable IPS events (related to end-user computing devices), the total TCO savings comes to **$37,125**. Again, most of these costs can be eliminated through proper use of automation.

Overall, without automation, our analysis shows that a network our size could conceivably achieve a TCO savings of **$184,845** through automation of tuning, assessment and user lookup, as summarized in Table 1.

Table 1. First-Year Savings Through Automation of Tuning, Assessment and User Lookup

Function	Costs Without Automation
IPS tuning	$ 39,720
Impact assessment	$ 108,000
Linking individual users with events	$ 37,125
Total first-year savings	**$ 184,845**

Reduced Exposure = Cost Avoidance

Let's face it. NGIPS should do its job in reducing or even eliminating exposures should an event occur that could result in a loss of data.

In this last part of our exercise, we consider the cost of a lost record. As noted in the Ponemon 2013 Cost of a Data Breach report,[4] the cost of a record lost to malicious attack is $275 per record. When an IPS captures an attack before it happens, it's hard to tell what the savings would be in terms of lost data. But we can examine some current cases in which data records were breached and estimate the cost avoidance that could be achieved for organizations with responsibility for personal data of value to attackers.

For that, let's take a look at the 2013 Verizon Data Breach Investigations Report.[5] In it, 66 percent of actual breaches investigated took months to discover, with 4 percent of those taking years to discover. In fact, the recent Mandiant report "2013 MTrends" determined that the median time for discovery of an attacker was 243 days in all the cases they studied.[6]

Immediate prevention before malicious code is executed and spread to other systems, of course, would be of ultimate value, but minimizing time to detection will also reduce costs of events that break past our defenses. As the Verizon report put it: "Without de-emphasizing prevention, focus on better and faster detection …."[7] The report continues, "Regularly measure things like 'number of compromised systems' and 'mean time to detection,' and use these numbers to drive better practices."

Calculations could also be used for determining TCO through a cost-avoidance model, when we consider the cost of losing records. In recent cases, organizations have lost from tens of thousands up to millions of records. Take, for example, the case of JPMorgan Chase & Co, which announced in 2013 that 465,000 cardholder accounts were breached by attackers that had made their way inside the Chase network.[8]

The attackers initially breached the network through its website in July and were not detected until September. Let's use this Chase breach to calculate the cost avoidance TCO:

Number of records breached:	**465,000**
The cost per hacked record:	**$275** ***(based on Verizon's analysis above)***
Overall cost of data loss:	**$127,875,000**

In addition to the overall costs, we must consider the incremental costs that accrue daily until the breach is detected and the losses are stopped. According to the Poneman Institute in their 2013 Cost of Cyber Crime Study,[9] an estimate of the daily losses until resolution of an "attack" averages $32,469 per day. Let's extend this over a 60-day period before the incident was discovered and data leakage was blocked. The final losses will increase on a daily basis until the "bleeding" is stopped. In this scenario, we could estimate an additional $1,948,140 lost due to the time delay (60 days x $32,469 per day).

---

[4] www.bankinfosecurity.com/interviews/data-breach-i-1953/op-1

[5] www.verizonenterprise.com/DBIR/2013

[6] www.mandiant.com/resources/mandiant-reports (requires registration)

[7] www.verizonenterprise.com/DBIR/2013, page 10

[8] www.reuters.com/article/2013/12/05/us-jpmorgan-dataexposed-idUSBRE9B405R20131205

[9] http://media.scmagazine.com/documents/54/2013_us_ccc_report_final_6-1_13455.pdf, page 13

Conclusion

The calculations in this document are designed to err on the side of being conservative. Even though conservative, these numbers reveal that there are true cost savings to be realized with the proper implementation of IPS automation features. Automation in tuning IPS policy, impact assessment and linking users with events are some of the newer capabilities in next-generation IPS toolsets that bring substantial savings. New automation capabilities now offer new synergies as NGIPS systems are learning to become more adaptive to their environments and policy requirements. NGIPS capabilities have been multiplied, for example, with the introduction of decryption tools that make formerly impenetrable network packets open, readable and actionable. Collected data provides added value to network forensic/analysis tools. In short, an NGIPS can add context to your network activity and everything connected to it, including the users.

A properly utilized NGIPS ultimately reduces TCO for an organization and helps mitigate risk against data loss to unwelcome malicious intruders. In this way, NGIPS prevents or reduces data losses that would otherwise have had a direct negative impact on the organization's bottom line.

Appendix A: A Case Study in Management TCO

One of the three largest credit reporting agencies implemented Sourcefire's NGIPS with automated impact assessment, user lookup and tuning. This multinational credit reporting organization has approximately 20,000 nodes and 7,500 total employees. The major motivation for installing the automated system was to greatly improve its security situational awareness through passive host fingerprinting.[10]

This company had considered bringing in a SIEM tool, but as a key security staff member for this company said, "SIEM is a very heavy lift for most companies." This person reported that the Sourcefire IPS tool can take in a wide variety of events and collect vulnerability data to approach the level and functionality of a SIEM tool without having to deal with the cost and operational impact of a SIEM tool.

Before Sourcefire's NGIPS product was installed, this organization had 20 Snort sensors that collected and sent a large volume of data. The situation became unmanageable because the sensors were unable to link and unify policy settings throughout the network. Furthermore, Snort does not fingerprint hosts. The Sourcefire IPS tool enabled this organization to integrate vulnerability data with operational security data, link and unify policy across the organization's entire enterprise, and tune policy settings as conditions and attacks changed. Furthermore, this tool enabled the organization to fingerprint hosts through passive fingerprinting, enabling it to determine which attacks were potentially able to succeed—and thus to greatly reduce the number of labor hours devoted to operational security monitoring.

While our source would not discuss actual dollars saved, he did discuss time saved, which we then calculated at the generic rate we set in our exercises to $75/hour. Table A-1 provides details of the calculations.

Table A-1. Summary of Calculated Savings

Function	Annual Savings	Explanation
IPS tuning	$54,000	It takes two weeks to manually tune policy (including shared policies), versus 2.5 days per month using automated tuning. Two weeks at 40 hours = $6,000 per month to manage policies manually. With automation, they're doing the same work in 20 hours per month, or $1,500 at $75/hour. That's a savings of $4,500 per month, or $54,000 per year.
Impact assessment	$119,700	Our source reported 160 man-hours per month manually analyzing the impact of events. At a cost of $75/hour, that equates to $12,000 a month to assess impact. With automation, the number of man-hours was reduced to one-sixth of that amount, or 27 hours per month, saving 133 hours per month ($9,975). Over 12 months, at $75/hour, that equates to a savings of $119,700 per year.
Linking individual users with events	$57,285	Approximately one-third (33%) of 200 actionable events per month are related to end-user systems configured for DHCP. At $75 an hour, the monthly expense of manually determining user identity for 67 events per month is $5,025. Now, this lookup is nearly instantaneous, reducing labor hours from an average of one hour down to three minutes per inquiry. So rather than $5,025 per month for 67 hours of work, it costs only $251.25 a month to look up users at three minutes per inquiry, saving $57,285 per year.
Overall annual savings	**$230,985**	The combination of automating IPS tuning, impact analysis and user identification results in a significant TCO cost reduction.

[10] Passive fingerprinting involves obtaining information about a network and the services and hosts therein by capturing data from traffic that flows through it. No active processes that alter the traffic and processes therein exist.

About the Authors

Eugene Schultz, Ph.D., CISM, CISSP, is CTO of Emagined Security and the author/coauthor of books on UNIX security, Internet security, Windows NT/2000 security, incident response, and intrusion detection and prevention. He was also the cofounder and original project manager of the Department of Energy's Computer Incident Advisory Capability (CIAC).

J. Michael Butler, GCFA, CISA, GSEC, EnCE, is an information security consultant with a leading provider of technical services for the mortgage industry. Butler's responsibilities have included computer forensics, information security policies (aligned to ISO and addressing federal and state disclosure laws), enterprise security incident management planning, internal auditing of information systems and infrastructure, service delivery and distributed systems support. He has also been involved in authoring SANS security training courseware, position papers, articles and blogs.

SANS would like to thank its sponsor:

SOURCEfire®

Last Updated: December 12th, 2017

Upcoming SANS Training

Click Here for a full list of all Upcoming SANS Events by Location

SANS Security East 2018	New Orleans, LAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS Amsterdam January 2018	Amsterdam, NL	Jan 15, 2018 - Jan 20, 2018	Live Event
SEC599: Defeat Advanced Adversaries	San Francisco, CAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
Northern VA Winter - Reston 2018	Reston, VAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
SANS Dubai 2018	Dubai, AE	Jan 27, 2018 - Feb 01, 2018	Live Event
SANS Las Vegas 2018	Las Vegas, NVUS	Jan 28, 2018 - Feb 02, 2018	Live Event
SANS Miami 2018	Miami, FLUS	Jan 29, 2018 - Feb 03, 2018	Live Event
Cyber Threat Intelligence Summit & Training 2018	Bethesda, MDUS	Jan 29, 2018 - Feb 05, 2018	Live Event
SANS London February 2018	London, GB	Feb 05, 2018 - Feb 10, 2018	Live Event
SANS Scottsdale 2018	Scottsdale, AZUS	Feb 05, 2018 - Feb 10, 2018	Live Event
SANS Southern California- Anaheim 2018	Anaheim, CAUS	Feb 12, 2018 - Feb 17, 2018	Live Event
SANS Secure India 2018	Bangalore, IN	Feb 12, 2018 - Feb 17, 2018	Live Event
Cloud Security Summit & Training 2018	San Diego, CAUS	Feb 19, 2018 - Feb 26, 2018	Live Event
SANS Dallas 2018	Dallas, TXUS	Feb 19, 2018 - Feb 24, 2018	Live Event
SANS Brussels February 2018	Brussels, BE	Feb 19, 2018 - Feb 24, 2018	Live Event
SANS Secure Japan 2018	Tokyo, JP	Feb 19, 2018 - Mar 03, 2018	Live Event
SANS New York City Winter 2018	New York, NYUS	Feb 26, 2018 - Mar 03, 2018	Live Event
CyberThreat Summit 2018	London, GB	Feb 27, 2018 - Feb 28, 2018	Live Event
SANS London March 2018	London, GB	Mar 05, 2018 - Mar 10, 2018	Live Event
SANS SEC460: Enterprise Threat Beta	OnlineCAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced